Efficient quantification of experimental evidence against local realism

Yanbao Zhang,[1, 2] Scott Glancy,[2] and Emanuel Knill[2]

[1]Department of Physics, University of Colorado Boulder, Boulder, Colorado, 80309, USA
[2]Applied and Computational Mathematics Division,
National Institute of Standards and Technology, Boulder, Colorado, 80305, USA

To quantify the evidence against local realism in an experiment, one can choose a test statistic to measure the amount of violation of local realism. It is desirable to bound the probability, according to local realism, of obtaining a test statistic at least as extreme as that observed. We describe an efficient protocol for computing such a bound from any set of Bell inequalities for any number of parties, measurement settings or outcomes. The bound depends on the choice and number of Bell inequalities, and more inequalities make the bound asymptotically tighter. We find that even trivial Bell inequalities such as those derived from no-signaling conditions can improve the tightness of the bound.

PACS numbers: 03.65.Ud, 03.65.Ta, 02.50.Tt

Theories designed according to "local realism" (LR) include a set of hidden variables, which if known would predict all measurement results; however, the values of the hidden variables cannot be influenced by spacelike-separated events. In 1964 Bell constructed an inequality satisfied by all correlations accessible by LR and showed that correlations between spacelike-separated measurements on two quantum systems can violate this inequality [1]. Since then, many experimental tests showing Bell-inequality violations have been performed (see Ref. [2] for a review). The importance of such a test is twofold. First, it shows that local realistic (LR) descriptions of bipartite quantum systems do not always exist. Second, it supports quantum information tasks such as quantum key distribution [3–5] and randomness generation [6, 7].

To test a Bell inequality in an experiment, one needs to estimate the probabilities of various outcomes from a finite number N of measurements. Due to uncertainties in the estimated probabilities, it is conventional to present the violation of LR in terms of the number of experimental standard deviations of violation of a Bell inequality. While this provides the precision with which a Bell-inequality violation is measured, it is not a valid indicator of the strength of experimental evidence against LR [8]. For the latter, one needs to take into account the possibility that N data points generated by an LR model can violate a Bell inequality due to statistical fluctuations in finite samples. This possibility can be formalized in statistics via a p-value for the hypothesis test of LR. We define a p-value as the maximum probability according to LR of obtaining a test statistic, such as a Bell-inequality violation, at least as extreme as that observed. Thus, a small p-value means that the observed data are significant for rejecting LR. Upper bounds of p-values for specified test statistics are required for precise statements on experimental evidence against LR. Such bounds not only help to reliably demonstrate violations of LR, but also help to prove the security of quantum key distribution or certify the generation of genuine randomness.

There are two available protocols that compute upper bounds of p-values. One is the martingale-based protocol [9, 10], but the bounds computed, such as the bound in Ref. [6], are not tight [8]. The other is the prediction-based-ratio (PBR) protocol [8], which computes tighter bounds. Specifically, the latter bounds are asymptotically tight with respect to N, if the prepared quantum states and measurement settings do not vary in time. While the PBR protocol is practical for many standard configurations, it is inefficient with respect to the number of parties per test, settings per party, and outcomes per setting. The reason is that it requires computing estimates of the experimental probability distribution and the associated optimal LR model. These estimates are difficult to find when there are many parties, settings, or outcomes. Extreme examples are provided by configurations involving continuous variables, where the PBR protocol cannot be directly applied. Here, we propose a simplified PBR protocol to efficiently compute high-quality p-value bounds for all configurations.

The simplified PBR protocol has at least three advantages over other protocols. First, the p-value bounds computed are as good as and typically better than those obtained by the martingale-based protocol. Second, multiple Bell inequalities can be taken into consideration at once in a statistically rigorous way. Thus we can obtain high-quality p-value bounds even when we cannot determine beforehand which inequality will work best. Third, this protocol can be applied to any test with linear witnesses, such as entanglement detection [11, 12], without a full analysis of the relevant probability space.

Preliminaries.—An experimental test of LR involves a number of trials. At each trial, each of a number of spatially separated parties performs a local measurement, where the setting is chosen randomly from a fixed set. Conventionally, at the end of the experiment, a predetermined Bell inequality is tested using the results from the trials. For example, if there are two parties and each party has two measurements with outcomes ± 1, the

Clauser-Horne-Shimony-Holt (CHSH) inequality [13]

$$E(A_1B_1) + E(A_1B_2) + E(A_2B_1) - E(A_2B_2) \le 2 \quad (1)$$

can be tested, where $E(A_iB_j)$ with $i,j \in \{1,2\}$ is the correlation between measurements A_i and B_j.

To statistically quantify the evidence against LR, we express a Bell inequality as an upper bound on the expectation of a function of a trial result. That is, we write a Bell inequality in the form $\langle I(X)\rangle \le B$, where I is a real-valued function, called a Bell function, and X is the random variable from which a trial result x is sampled. The result x consists of the measurement-setting choices made by all parties and the outcomes of these measurements. To write a Bell inequality in the above form, the probability distribution of the joint measurement-settings is assumed to be known and fixed before an experiment. There is no loss of generality in assuming this, as explained in Refs. [8–10]. For example, for the CHSH inequality (1), a trial result x consists of setting choices i,j and outcomes a_i, b_j, and we can write $I_{\mathrm{CHSH}}(x) = 4(1-2\delta_{i,2}\delta_{j,2})a_ib_j$ and $B = 2$, where we have assumed that the joint-setting distribution is uniform.

As explained in the introduction, we quantify the strength of experimental evidence against LR by means of a p-value. A p-value is associated with a test statistic T that is a function of the sequence of trial results. If N is the total number of trials, the corresponding sequence of results is denoted by $\mathbf{x} = (x_1, \ldots, x_N)$. As is conventional, we distinguish between the sequence of results and the sequence of random variables $\mathbf{X} = (X_1, \ldots, X_N)$ giving rise to these results. The exact p-value p_N is defined as the maximum of the probabilities of the events $T(\mathbf{X}_{\mathrm{LR}}) \ge T(\mathbf{x})$ over all random-variable sequences $\mathbf{X}_{\mathrm{LR}}$ distributed according to LR models. That is

$$p_N = \max_{\mathrm{LR}} \mathrm{Prob}_{\mathrm{LR}}(T(\mathbf{X}_{\mathrm{LR}}) \ge T(\mathbf{x})). \quad (2)$$

Due to the difficulty of determining worst-case tail probabilities of typical test statistics, we can usually determine only upper bounds of exact p-values. Thus, for the remainder of the paper, the term "p-value" refers to any valid upper-bound on the exact p-value. For the protocols discussed below, if the trial results are independent and identically distributed according to a distribution that violates LR, the p-values computed decrease to 0 exponentially as $N \to \infty$. We can therefore compare different protocols' performances in a test of LR according to the confidence-gain rate defined by

$$G = -\lim_{N\to\infty} \frac{\log_2 p_N^{(\mathrm{prot})}}{N}, \quad (3)$$

where $p_N^{(\mathrm{prot})}$ is the p-value computed by a protocol. Higher gain rates imply better protocol performance. Each protocol discussed below works even under memory effects [14], that is, even when the prepared quantum state, measurement settings, and relevant LR models vary arbitrarily with time.

PBR protocols.—The test statistic used by a PBR protocol is based on nonnegative functions R_n to be applied to the n'th trial result x_n and satisfying $\langle R_n(X)\rangle \le 1$ for X distributed according to any LR model. Thus, each R_n is a nonnegative Bell function for the Bell inequality $\langle R_n(X)\rangle \le 1$. The function R_n is constructed *before* observing the n'th trial result x_n. Its construction can use information from previous trials and typically requires predicting the distribution of X_n. Thus, R_n is referred to as a prediction-based ratio (PBR). A PBR protocol computes a test statistic according to $T(\mathbf{x}) = \prod_{n=1}^{N} R_n(x_n)$. To obtain a p-value for $T(\mathbf{x})$ it suffices to observe that by construction T is nonnegative and $\langle T(\mathbf{X}_{\mathrm{LR}})\rangle \le 1$, so that by Markov's inequality we can compute a p-value according to

$$p_N^{(\mathrm{PBR})} = \min(1/T(\mathbf{x}), 1). \quad (4)$$

See Ref. [8] for further details. Different PBR protocols are characterized by how they choose the PBR R_n for each n.

Full PBR protocol.—For the full PBR protocol [8], R_n is chosen so as to optimize the expected confidence-gain rate given previous trial results. For this optimization, the protocol assumes that X_n's distribution is the same as that from which the trial results $x_1, \ldots, x_{n-1}$ were sampled, and that these samples are independent. Whether or not these assumptions actually hold affects only the quality of the p-value computed, but not its validity. Given these assumptions, R_n is computed in two steps. The first is to make an estimate of the experimental probability distribution $q(x) \equiv \mathrm{Prob}_{\mathrm{QM}}(X_n = x)$, and the second is to determine the probability distribution $p(x) \equiv \mathrm{Prob}_{\mathrm{LR}}(X_{\mathrm{LR}} = x)$ according to the LR model that minimizes the Kullback-Leibler (KL) divergence from the estimate q [15]. The protocol then sets the next PBR to $R_n(x_n) = q(x_n)/p(x_n)$. Details about this protocol and the proof that this R_n satisfies the conditions on a PBR are in Ref. [8].

While the implementation of the full PBR protocol is practical for typical experimental tests of LR, its complexity is not well-behaved as the number of parties, settings, or outcomes increases. In particular, on each PBR computation, it needs to optimize a convex objective function over the convex space of all LR models, where the evaluation of the objective function involves a sum over all possible setting and outcome combinations of the parties. In contrast, the simplified PBR protocol introduced next requires an optimization over a small-dimensional convex space of a convex objective function that can be written as a sum of at most $(N-1)$ terms.

Simplified PBR protocol.—The simplified PBR protocol chooses the PBRs from the convex combinations of

Bell functions that are derived from a given set of Bell inequalities. To ensure that a convex combination is a PBR, the Bell functions first need to be standardized so that they are nonnegative and have expectations at most 1 for any LR model. Any Bell function that is lower-bounded has such a standardized form. In particular, if $\langle I(X)\rangle \leq B$ is a Bell inequality and $I(x) \geq b$ for all x, then $r(x) = (I(x) - b)/(B - b)$ is standardized. Note that, as a constraint on the distribution of X, $\langle r(X)\rangle \leq 1$ is equivalent to $\langle I(X)\rangle \leq B$. Given Bell inequalities $\langle I^{(m)}(X)\rangle \leq B^{(m)}$ where $I^{(m)}$ is lower-bounded and $m = 1, 2, \ldots, M$, we can construct the corresponding standardized Bell functions $r^{(m)}$. We define $\mathbf{r} = (r^{(1)}, \ldots, r^{(M)})$. The simplified PBR protocol chooses the PBR R_n from among the convex combinations

$$\boldsymbol{\omega} \cdot \mathbf{r} = \sum_m \omega_m r^{(m)}, \tag{5}$$

where $\omega_m \geq 0$ and $\sum_m \omega_m = 1$. Our implementation always includes the trivial Bell function $r^{(1)} = 1$. This ensures that the set of convex combinations is at least one-dimensional and that the confidence-gain rate is at least as high as that achieved by the martingale protocol.

Like the full PBR protocol, the simplified PBR protocol aims to optimize the expected confidence-gain rate given previous trial results, under the assumption that the distribution of X_n is the same as the empirical-frequency distribution of the previous trial results. Whether or not this assumption holds does not affect the validity of the p-value computed. The confidence gain of the n'th trial may be defined as $\log_2 R_n(x_n)$. Its expected value given that X_n is distributed according to q is

$$\sum_{x_n} q(x_n) \log_2 R_n(x_n). \tag{6}$$

Before the n'th trial, the protocol attempts to maximize this expected confidence gain. Since q is not known, it is empirically estimated based on the first $(n-1)$ trials. Expanding R_n according to Eq. (5) yields the following estimate of the expected confidence gain of the n'th trial:

$$G_n(\boldsymbol{\omega}) = \frac{1}{n-1} \sum_{k=1}^{n-1} \log_2(\boldsymbol{\omega} \cdot \mathbf{r}(x_k)). \tag{7}$$

The protocol thus determines R_n by maximizing $G_n(\boldsymbol{\omega})$ over $\boldsymbol{\omega}$, that is, $R_n = \mathbf{r} \cdot \mathrm{argmax}_{\boldsymbol{\omega}} G_n(\boldsymbol{\omega})$. Note that, unlike the full PBR protocol, the simplified PBR protocol does not require explicitly optimizing over all LR models. Computing $\mathrm{argmax}_{\boldsymbol{\omega}} G_n(\boldsymbol{\omega})$ requires optimizing a convex objective function over an M-dimensional convex space, where the evaluation of the objective function involves a sum of $(n-1)$ terms. In our implementation, we apply the expectation-maximization (EM) algorithm [16] to solve this problem.

The performance of the simplified PBR protocol depends on the relationship between the actual distribution of trial results and the set of standardized Bell functions used. If the results are independent and identically distributed according to a known distribution that violates LR, then there exists an optimal Bell inequality that can be derived from the optimal PBR as found by the full PBR protocol. (Here, optimality refers to the optimality of the gain rate achieved by the protocol; see Ref. [8].) If the optimal PBR is included in the convex set of standardized Bell functions, the confidence-gain rate achieved by the simplified PBR protocol is optimal. But since the actual distribution is unknown before an experiment, the above assumption may not hold without making the dimension of the set of convex combinations in Eq. (5) impractically large. Thus, before an experiment, it is important to choose a relevant (and preferably small) set of standardized Bell functions. Below we show that it helps to include more than just the obvious Bell functions.

The performance of the simplified PBR protocol can be compared with that of the martingale-based protocol [9, 10], the only valid non-PBR protocol considered so far. The martingale-based protocol uses a Bell inequality $\langle I(X)\rangle \leq B$ with a bounded Bell function, chosen before an experiment. After the experiment, the mean of the Bell function is estimated as $\hat{I} = \frac{1}{N}\sum_{n=1}^{N} I(x_n)$. To obtain a p-value, the protocol uses $\hat{I}$ as the test statistic. If $\hat{I} \geq B$, the bounds on the Bell function imply that a p-value can be computed according to

$$p_N^{(\mathrm{mart})} = \left[\left(\frac{a-B}{a-\hat{I}}\right)^{\frac{a-\hat{I}}{a-b}} \left(\frac{B-b}{\hat{I}-b}\right)^{\frac{\hat{I}-b}{a-b}}\right]^N, \tag{8}$$

where $a = \sup_x I(x)$ and $b = \inf_x I(x)$. This p-value expression is based on a version of Hoeffding's bound in Ref. [17], which improves the ones given in Refs. [6, 8, 10]. The derivation of Eq. (8) is explained in the supplemental material (SM). In the SM, we also show that the simplified PBR protocol using the same Bell inequality, together with the default trivial Bell function $r = 1$, achieves a gain rate at least as high as the gain rate achieved by the martingale-based protocol. These two gain rates are equal to each other if and only if the experimental range of the function I is contained in the set $\{a, b\}$.

Computational resource costs.—Of the available protocols for computing p-values, the martingale-based one is the least resource-intensive and simplest to apply. It requires computing only an estimate of the mean of the Bell function, which involves a sum of N terms. The motivation for applying a PBR protocol is that it can adapt to changes in the experimental results' distribution and find better Bell inequalities with respect to the

observed data, both of which are difficult to do before an experiment.

For quantifying the resource cost of the simplified PBR protocol, we assume that the Bell functions used can be evaluated in constant time given the values of the arguments. This assumption is realistic for many Bell functions, as their values are determined by concise formulas derived from theory. Alternatively, these functions can be preprocessed as a table stored in random-access memory; we do not include preprocessing time in our analysis. We also assume that an optimization is performed by algorithms whose resource costs are determined primarily by the complexity C of evaluating the objective function and the dimension D of the convex search space. In particular, we do not account for algorithm-specific resource costs such as the number of iterations required to achieve sufficient precision for our purpose. Also, we do not account for the complexity of enforcing convex constraints. This is motivated by the observation that there is no additional overhead for enforcing convex constraints in the EM algorithm used in our implementation. Given our assumptions and from Eq. (7), the complexity of evaluating the objective function in the optimization required for constructing the PBR R_n is $C = O(nM)$, which is the product of the number of terms in the sum and the number of Bell-function evaluations underneath the logarithm. The dimension of the search space is $D = O(M)$, the size of $\boldsymbol{\omega}$ in Eq. (7). Consequently, unlike the full PBR protocol (see the details in the SM), the numbers of parties, settings, and outcomes are not limiting factors. In this sense, the simplified PBR protocol is efficient for any experimental configuration.

In our discussion so far, we have assumed that each PBR protocol updates the PBR before each trial. In practice, the PBR is updated only for a block of trial results at a time (see Ref. [8]), thus limiting PBR computations to when enough new information has been obtained, thereby reducing the resource cost.

Protocol comparison.—We begin by comparing the confidence-gain rates achieved by different protocols for experimental configurations designed to violate the Collins-Gisin-Linden-Massar-Popescu (CGLMP) inequality [18]. To test the CGLMP inequality, there are two parties, and each of them performs one of two possible measurements with d outcomes at each trial. This is an example where the full PBR protocol is impractical for large d. For this example and the one below, we assume that at each trial each party's measurement setting is chosen uniformly randomly. The CGLMP inequality can be written as $\langle I_d(X)\rangle \leq 2$, where the function I_d takes d different values. The gain rates G_{mart} and G_{sPBR}, achieved by the martingale-based and simplified PBR protocols, are shown in Fig. 1. Here the simplified PBR protocol uses only the CGLMP inequality. This figure illustrates that G_{sPBR} is higher than G_{mart} when $d > 2$.

The optimal gain rate S_q is achieved by the full PBR protocol and can be computed as the minimum KL divergence from the experimental probability distribution to any LR model [20]. For the results of Fig. 1, we find that the gain rates G_{sPBR} are numerically indistinguishable from S_q when $d \leq 13$. For the case $d > 13$, it is difficult to compute S_q due to the large dimension of the probability space over all possible LR models. For the tests studied in Fig. 1, we conjecture that $G_{\mathrm{sPBR}} = S_q$. In general we cannot guarantee that G_{sPBR} is optimal.

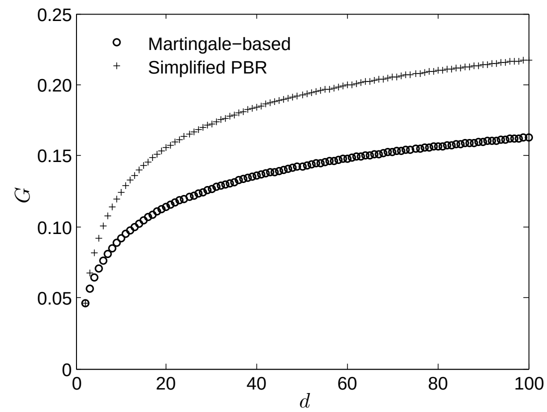

FIG. 1: Confidence-gain rates in the test of the CGLMP inequality $\langle I_d(X)\rangle \leq 2$. Here, we use the quantum state and measurement settings of Ref. [19], Eqs. (15) and (9), respectively.

Next, we compare the performance of the simplified PBR protocol when using different numbers of Bell inequalities. The experimental configuration considered is for a test of the CHSH inequality (1) using an unbalanced Bell state $|\psi(\theta)\rangle = \cos(\theta)|00\rangle + \sin(\theta)|11\rangle$. For comparison, we consider the simplified PBR protocol with the CHSH inequality alone or in conjunction with additional, seemingly trivial Bell inequalities such as those derived from no-signaling conditions. With Bell functions corresponding to no-signaling conditions, the gain rates are improved, as shown in Fig. 2.

In the SM, we show how the p-values computed by different protocols behave in a simulated experiment as functions of the number of trials.

Extensions.—To compute a p-value, the simplified PBR protocol uses a set of linear inequalities that are satisfied by the predictions of a null hypothesis before each trial in an experiment. Besides tests of LR, there are many other types of tests based on linear witnesses, such as tests for entanglement [11, 12] and system dimensionality above a given bound [21, 22]. In any test based on linear witnesses, such a witness can be expressed as $\langle W(X)\rangle \leq B$, where W is a real-valued function and X is the random variable from which a trial result x is sampled. The result x consists of all choices made at each trial, such as choices of states and measurement settings, and the outcomes observed under these choices. Here,

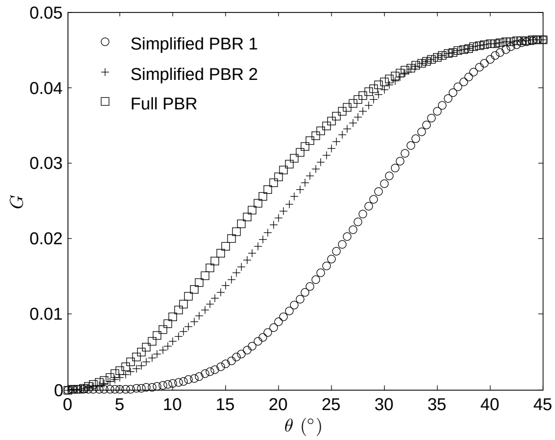

FIG. 2: Confidence-gain rates in the test of LR with an unbalanced Bell state $|\psi(\theta)\rangle$. The measurement settings are chosen to maximize the violation of the CHSH inequality (1) given the state $|\psi(\theta)\rangle$. The gain rates achieved by the simplified PBR protocol using the CHSH inequality are shown as circles ($\circ$), while the gain rates by the same protocol using the CHSH inequality together with no-signaling conditions are shown as crosses ($+$).

we assume that the choices are made randomly according to a known probability distribution at each trial, so that a witness $\langle W(X)\rangle \leq B$ is satisfied before each trial assuming the null hypothesis. As for Bell functions, if a witness W is lower-bounded it can be standardized. The simplified PBR protocol can then be applied with any set of standardized witnesses, as we did in a test of LR.

We thank Kevin Coakley for helpful comments and discussions. This paper is a contribution of the National Institute of Standards and Technology and is not subject to U.S. copyright.

Supplemental material

Computational resource comparison

In this section, we compare the computational resources required by the simplified and full PBR protocols in an experimental test of LR.

We consider an experimental configuration involving l parties where each party has s measurement settings and each local measurement has d outcomes. (The comparison below is readily extended to more general configurations.) We suppose that the joint-setting distribution is uniform. Then, the number of possible results (measurement settings and outcomes of all parties) at a trial is $K = (ds)^l$. Since a deterministic LR model specifies the exact outcome for each local measurement of each party at a trial, there are $H = d^{ls}$ many such models. A general LR model is a convex combination of deterministic LR models, so the number of free parameters characterizing a general LR model is $(H - 1)$.

Let the total number of trials in an experimental test of LR be N. We assume that each PBR protocol sets the initial value of the PBR to $R_1 = 1$ and updates the PBR R_n before each trial n ($n > 1$). (In practice this is unnecessary; see the appendix of Ref. [8].) For updating the PBR, each PBR protocol needs to optimize a convex objective function over a convex space. The complexity of this optimization problem can be described in terms of variables that are functions of the parameters n, l, s, and d characterizing the input data size. (Note that the stored size of the first n trial results is $O(n\log(K)) = O(nl(\log(d) + \log(s)))$.) We need to quantify the resource cost of implementing each protocol in terms of these parameters.

The complexity of the optimization problem solved before each trial can be parametrized by the complexity of the convex search space, the complexity of evaluating the objective function, and the precision needed for computing a high-quality p-value for rejecting LR. We assume that the simplified and full PBR protocols use generic iterative optimization algorithms whose implementation complexities as functions of these parameters are asymptotically the same. We also assume that the complexity of the convex search space is dominated by its dimension. In particular, we do not account for the complexity of enforcing convex constraints. For quantifying the complexity of evaluating the objective function, we assume that the Bell functions used can be evaluated in constant time given any trial result, as explained in the main text. Also, we assume that determining whether or not an arbitrary trial result x happens according to a deterministic LR model takes constant time. (Strictly speaking, the time taken for such a determination process is proportional to the number of parties l.) The precision needed affects the number of iterations required by an algorithm to find a numerical solution. It affects only the quality of the p-value computed by a protocol, but not its validity. (For the EM algorithm [16] used, see Theorem 4 of Ref. [16] and the appendix of Ref. [8] for the effects of the precision parameters in the simplified and full PBR protocols, respectively.) We assume that the precision parameters in both protocols are set to be the same, and we do not account for the number of iterations required to achieve the specified precision. Therefore, for the purpose of comparing the computational resources required by the simplified and full PBR protocols, we focus on comparing the dimensions D of the convex search spaces and the complexities C of evaluating the objective functions in the optimization problems solved by the two protocols before each trial.

We first consider the simplified PBR protocol. Given a set of M Bell inequalities, this protocol sets $R_n = \boldsymbol{\omega}_n \cdot \mathbf{r}$, where the size of $\boldsymbol{\omega}_n$ is M, $\mathbf{r}$ is defined before Eq. (5) in the main text, and $\boldsymbol{\omega}_n$ is chosen to maximize the esti-

mated confidence gain

$$\frac{1}{n-1}\sum_{k=1}^{n-1}\log_2(\boldsymbol{\omega}\cdot\mathbf{r}(x_k)) = \sum_{x:f_n(x)\neq 0} f_n(x)\log_2(\boldsymbol{\omega}\cdot\mathbf{r}(x)), \tag{A.9}$$

where $f_n(x)$ is the empirical frequency of x before the n'th trial. Note that, in the right-hand side of Eq. (A.9), the sum is taken over only the results x already observed in the previous trials.

For the maximization of Eq. (A.9), the dimension of the convex search space is M. The evaluation of the objective function can use the left-hand or right-hand side of Eq. (A.9), whichever has fewer terms. Thus it involves a sum of at most $\min(n-1, K)$ terms where each term requires computing a convex combination of M Bell-function values. Hence, for updating the PBR R_n before the n'th trial, the complexity of evaluating the objective function is $C_{\text{sPBR}} = O(\min(nM, KM)) = O(\min(nM, (ds)^l M))$, and the dimension of the search space is $D_{\text{sPBR}} = O(M)$. Therefore, if any of the configuration parameters l, s, or d is large, C_{sPBR} and D_{sPBR} are independent of these parameters, and so the simplified PBR protocol can be applied efficiently.

The full PBR protocol [8] computes R_n in two steps. First, the protocol estimates the probability $q(x)$ of the result x to be observed at the next trial. This estimate can be obtained in different ways. The simplest is to let $q(x)$ be the empirical frequency $f_n(x)$ of x over the previous $(n-1)$ trials. However, one can consider additional constraints such as the known joint-setting distribution and no-signaling conditions. Thus, in Ref. [8] we suggested maximizing the log-likelihood function $L(q') \propto \sum_x f_n(x)\log_2(q'(x))$, subject to these constraints, and we observed that this can improve the quality of the p-value computed. Since this maximization is not a resource bottleneck, we do not consider its complexity in the comparison. Second, we find the LR model p closest to the estimated distribution q by minimizing the KL divergence [15] from q to an LR model p_{LR}

$$D_{\text{KL}}(q|p_{\text{LR}}) = \sum_x q(x)\log_2\frac{q(x)}{p_{\text{LR}}(x)}. \tag{A.10}$$

The full PBR protocol then sets $R_n(x_n) = q(x_n)/p(x_n)$.

For the minimization of Eq. (A.10), the dimension of the convex search space is H. The evaluation of the objective function involves a sum of K terms where each term requires computing $p_{\text{LR}}(x)$ according to a convex combination of H deterministic LR models. Hence, for updating the PBR R_n before the n'th trial, the complexity of evaluating the objective function is $C_{\text{fPBR}} = O(KH) = O(d^{l(s+1)}s^l)$, and the dimension of the search space is $D_{\text{fPBR}} = O(H) = O(d^{ls})$. While C_{fPBR} and D_{fPBR} are polynomial in d, they are exponential in each of l and s. Therefore, the full PBR protocol is not efficient with respect to these configuration parameters.

Before applying the simplified PBR protocol, one chooses a relevant and preferably small set of Bell inequalities. In many cases of interest, l, s, or d is large, and so is $H = d^{sl}$. For example, in field-quadrature measurements d is fundamentally infinite. Hence, M, the number of Bell inequalities used in the simplified PBR protocol, is in general much smaller than H, the number of deterministic LR models considered in the full PBR protocol. The complexities show that for such cases, the simplified PBR protocol is substantially less resource-intensive than the full PBR protocol.

The martingale-based protocol's p-value

Consider a Bell inequality $\langle I(X)\rangle \leq B$ with a Bell function I whose range is included in the interval $[b, a]$, where $b \leq a$. An experimental test yields an estimate $\hat{I} = \frac{1}{N}\sum_{n=1}^{N} I(x_n)$ of the mean of I, where $x_1, \ldots, x_N$ are the trial results.

Suppose that the n'th trial result x_n is distributed according to a random variable $X_{\text{LR},n}$ satisfying LR. In this case, the random variable from which $\hat{I}$ is sampled is $I_{\text{LR}} = \frac{1}{N}\sum_{n=1}^{N} I(X_{\text{LR},n})$. The sequence $M_n = \sum_{k=1}^{n}(I(X_{\text{LR},k}) - B)$ is a super-martingale, as shown in Refs. [9, 10]. Thus, for $t \geq 0$, the probability

$$\text{Prob}_{\text{LR}}\left(M_N \geq Nt\right) \leq \left[\left(\frac{a-B}{a-B-t}\right)^{\frac{a-B-t}{a-b}}\left(\frac{B-b}{B+t-b}\right)^{\frac{B+t-b}{a-b}}\right]^N. \tag{A.11}$$

The inequality (A.11) follows from Theorem 6.1 of Ref. [23]. Since $\text{Prob}_{\text{LR}}(I_{\text{LR}} \geq \hat{I}) = \text{Prob}_{\text{LR}}(M_N \geq N(\hat{I} - B))$, from the above inequality (A.11) we get the p-value of Eq. (8) in the main text.

Note that, although Theorem 6.1 of Ref. [23] is stated for a martingale, the same result and its proof also apply to a super-martingale. The same bound is also derived in Theorem 1 of Ref. [17] for a sum of independent random variables. From Refs. [17, 23], we can see that the bound in Eq. (A.11) is tighter than bounds of $\text{Prob}_{\text{LR}}\left(M_N \geq Nt\right)$ used in previous works [6, 8, 10] and derived from Azuma's inequality [23, 24].

Proof of $G_{\textbf{mart}} \leq G_{\textbf{sPBR}}$

We suppose that the martingale-based protocol uses a Bell inequality $\langle I(X)\rangle \leq B$ with a bounded Bell function such that $b \leq I(x) \leq a$ for all x. Also, we suppose that the simplified PBR protocol uses the standardized form of this Bell inequality together with the trivial Bell function $r = 1$.

Let the experimental probability of observing the result x be $q(x)$. The experimental mean of I is $I_q = \int q(x)I(x)\mathrm{d}x$. If $I_q \geq B$, then from Eqs. (3) and (8) in the main text we get the gain rate

$$\begin{aligned} G_{\mathrm{mart}} &= \frac{a-I_q}{a-b}\log_2\frac{a-I_q}{a-B} + \frac{I_q-b}{a-b}\log_2\frac{I_q-b}{B-b} \\ &= \int q(x)\left(\frac{a-I(x)}{a-b}\log_2\frac{a-I_q}{a-B} + \frac{I(x)-b}{a-b}\log_2\frac{I_q-b}{B-b}\right)\mathrm{d}x. \end{aligned} \tag{A.12}$$

Here, we use the fact that the experimental estimate $\hat{I}$ approaches I_q as $N \to \infty$. By the concavity of $\log_2(x)$ and some algebra, we get that the gain rate G_{mart} satisfies the inequality

$$\begin{aligned} G_{\mathrm{mart}} &\leq \int q(x)\log_2\left(\frac{a-I(x)}{a-b}\frac{a-I_q}{a-B} + \frac{I(x)-b}{a-b}\frac{I_q-b}{B-b}\right)\mathrm{d}x \\ &= \int q(x)\log_2\left(\omega_0\frac{I(x)-b}{B-b} + 1-\omega_0\right)\mathrm{d}x, \end{aligned} \tag{A.13}$$

where $0 \leq \omega_0 = \frac{I_q-B}{a-B} \leq 1$.

From Eqs. (3) and (4) in the main text and according to the design of the PBRs by the simplified PBR protocol (as explained in the main text), the gain rate achieved by this protocol is

$$G_{\mathrm{sPBR}} = \max_{0\leq\omega\leq 1}\int q(x)\log_2\left(\omega\frac{I(x)-b}{B-b} + 1-\omega\right)\mathrm{d}x. \tag{A.14}$$

Here, we use the fact that the empirical frequency $f_N(x)$ approaches the experimental probability $q(x)$ as $N \to \infty$. The inequality $G_{\mathrm{mart}} \leq G_{\mathrm{sPBR}}$ follows from comparing Eq. (A.13) with Eq. (A.14).

By considering the condition for equality in Eq. (A.13), we can show that $G_{\mathrm{mart}} = G_{\mathrm{sPBR}}$ if and only if $q(x) = 0$ whenever $b < I(x) < a$. For this it suffices to note that $\log_2(x)$ is strictly concave, so equality holds in Eq. (A.13) if and only if $I(x) = a$ or b whenever $q(x) \neq 0$.

Behavior of the protocols for finite data

Here we consider the behavior of each protocol given a finite amount of experimental data. We simulate the test of the CGLMP inequality $\langle I_3(X)\rangle \leq 2$ [18] with the quantum state and measurement settings of Ref. [19], Eqs. (15) and (9) (with $d = 3$), respectively. We assume that at each trial each party's measurement setting is chosen uniformly randomly. The protocols' gain rates are $G_{\mathrm{mart}} = 0.0565$ and $G_{\mathrm{sPBR}} = 0.0675$, while the optimal gain rate S_q achieved by the full PBR protocol is numerically indistinguishable from G_{sPBR}. For computing G_{sPBR}, the simplified PBR protocol uses the standardized CGLMP inequality and the trivial Bell function $r = 1$.

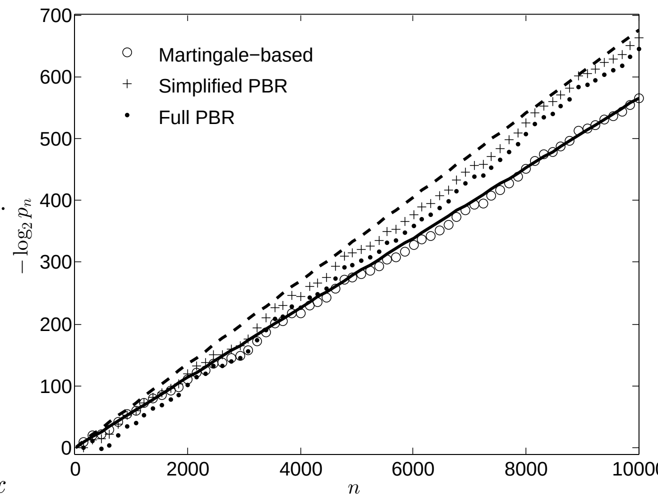

FIG. 3: An example of running log-p-values as functions of the number of trials n in a test of the CGLMP inequality. The dashed and solid lines are the asymptotic lines for log-p-values based on gain rates according to the (full or simplified) PBR protocol and the martingale-based protocol, respectively. Repetitions of this Monte Carlo simulation show similar behavior.

The results from 10,000 successive trials are recorded. Fig. 3 shows the (negative) log-p-values computed for the first n results from a simulated sequence of trials as functions of n. The asymptotic lines for log-p-values, given by the products of n and the respective gain rates achieved by different protocols, are also shown in Fig. 3.

For the simulation shown in Fig. 3, we update the PBRs and log-p-values only after every block including 154 successive trials. (See our previous work [8] for a discussion of the block-size choice and related issues.) This improves the efficiency of a PBR protocol. It also mitigates the offset of the computed log-p-values from the asymptotic line. This offset is due to an initial transient where the relevant features of the experimental distribution are being learned. The learning offset can be removed if, before an experiment, we have a good estimate of the experimental results' distribution. Such an estimate could be based on (quantum or otherwise) theory or previous experiments.

The PBR protocols provide better results than the martingale-based protocol. However, the PBR log-p-values show learning offsets from the asymptotic line. Our results show that the simplified PBR log-p-values have a smaller learning offset than the full PBR log-p-values in each of 30 independent simulations performed. The reason is that the simplified PBR protocol needs to infer a much smaller number of parameters for constructing the PBRs.

In the above example, the simplified PBR protocol uses only two Bell functions. Given a prescient choice of Bell functions, this is sufficient for computing asymptotically

optimal p-values. But in general, more Bell functions are needed for computing a high-quality p-value. However, this involves inferring more parameters and thus requires more trials before a good inference can be obtained. As a result, the learning offset is expected to increase when using more Bell functions. One way to mitigate this problem may be to increase the number of Bell functions used over time, adding new Bell functions only when there are enough trials for reliable inference of the additional parameters.

[1] J. S. Bell, Physics **1**, 195 (1964).
[2] M. Genovese, Phys. Rep. **413**, 319 (2005).
[3] J. Barrett, L. Hardy, and A. Kent, Phys. Rev. Lett. **95**, 010503 (2005).
[4] L. Masanes, Phys. Rev. Lett. **102**, 140501 (2009).
[5] L. Masanes, S. Pironio, and A. Acin, Nat. Commun. **2**, 238 (2011).
[6] S. Pironio *et al.*, Nature **464**, 1021 (2010).
[7] R. Colbeck and A. Kent, J. Phys. A: Math. Theor. **44**, 095305 (2011).
[8] Y. Zhang, S. Glancy, and E. Knill, Phys. Rev. A **84**, 062118 (2011).
[9] R. D. Gill, in *Mathematical Statistics and Applications: Festschrift for Constance van Eeden. Eds: M. Moore, S. Froda and C. Léger. IMS Lecture Notes – Monograph Series* (Institute of Mathematical Statistics. Beachwood, Ohio, 2003), vol. 42, pp. 133–154, also available as arXiv:quant-ph/0110137.
[10] R. D. Gill, in *Proc. of "Foundations of Probability and Physics - 2", Ser. Math. Modelling in Phys., Engin., and Cogn. Sc.* (Växjö Univ. Press., 2003), vol. 5, pp. 179–206, also available as arXiv:quant-ph/0301059.
[11] M. Horodecki, P. Horodecki, and R. Horodecki, Phys. Lett. A **223**, 1 (1996).
[12] B. M. Terhal, Phys. Lett. A **271**, 319 (2000).
[13] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, Phys. Rev. Lett. **23**, 880 (1969).
[14] J. Barrett, D. Collins, L. Hardy, A. Kent, and S. Popescu, Phys. Rev. A **66**, 042111 (2002).
[15] S. Kullback and R. A. Leibler, Ann. Math. Statist. **22**, 79 (1951).
[16] T. M. Cover, IEEE Trans. Inform. Theory **30**, 369 (1984).
[17] W. Hoeffding, Journ. Amer. Statist. Assoc. **58**, 13 (1963).
[18] D. Collins, N. Gisin, N. Linden, S. Massar, and S. Popescu, Phys. Rev. Lett. **88**, 040404 (2002).
[19] J.-L. Chen, C. Wu, L. C. Kwek, C. H. Oh, and M.-L. Ge, Phys. Rev. A **74**, 032106 (2006).
[20] R. R. Bahadur, in *Proc. Fifth Berkeley Symp. on Math. Statist. and Prob.* (Univ. of Calif. Press, 1967), vol. 1, pp. 13–26.
[21] N. Brunner, S. Pironio, A. Acin, N. Gisin, A. A. Méthot, and V. Scarani, Phys. Rev. Lett. **100**, 210503 (2008).
[22] R. Gallego, N. Brunner, C. Hadley, and A. Acín, Phys. Rev. Lett. **105**, 230501 (2010).
[23] C. McDiarmid, in *Surveys in Combinatorics* (Cambridge Univ. Press, 1989), vol. 141 of *London Math. Soc. Lecture Notes*, pp. 148–188.
[24] K. Azuma, TohoKu Math. Journ. **19**, 357 (1967).